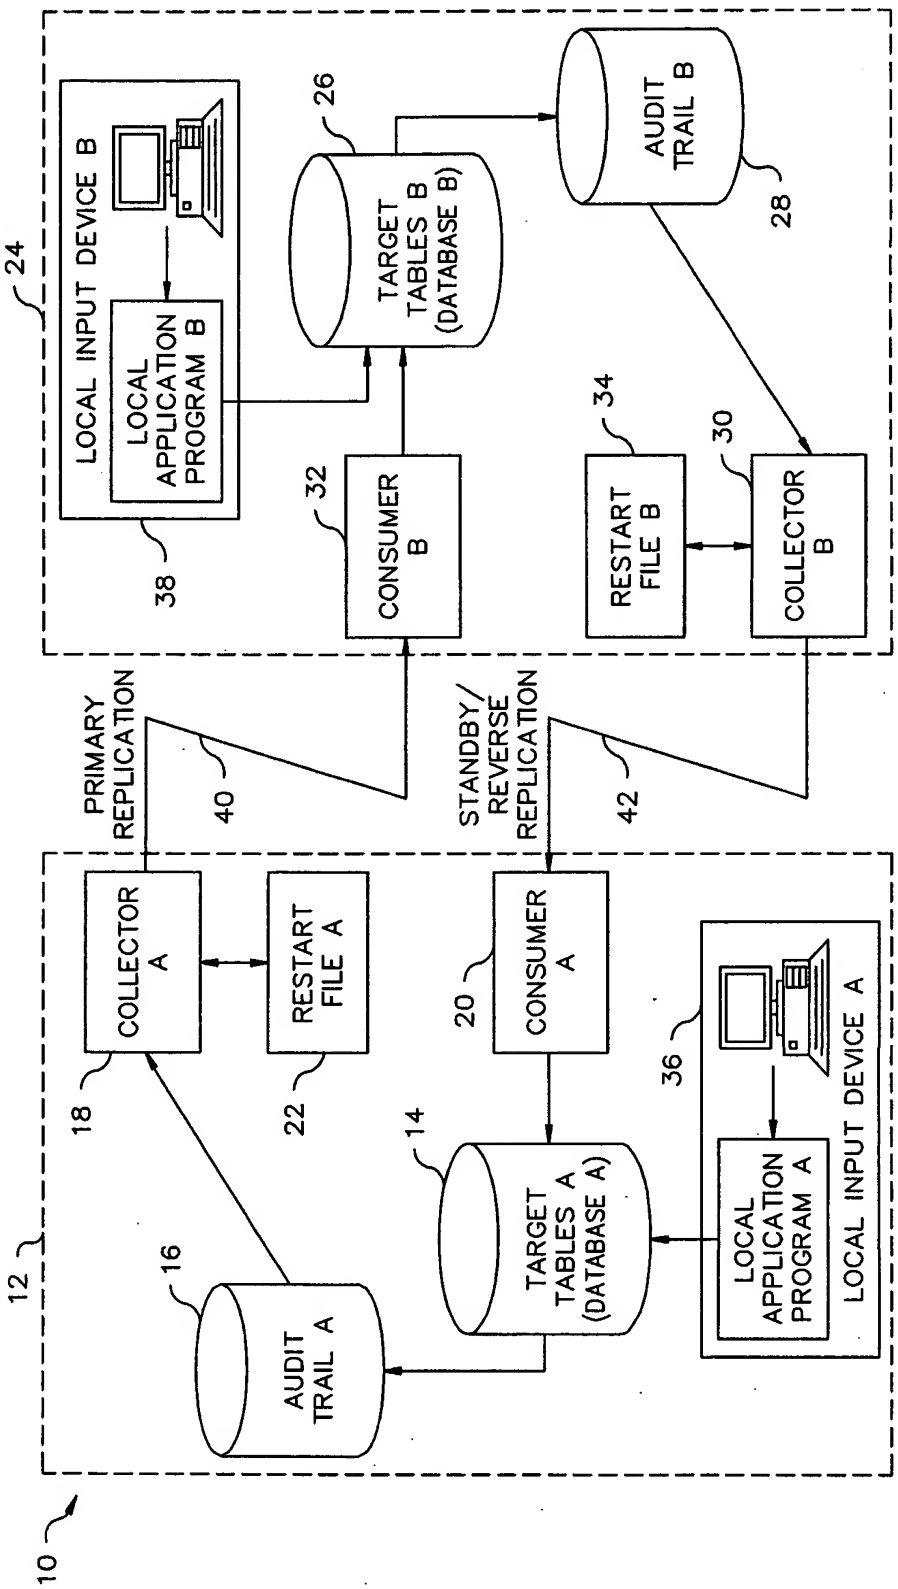

Fig. 1
(Prior Art)

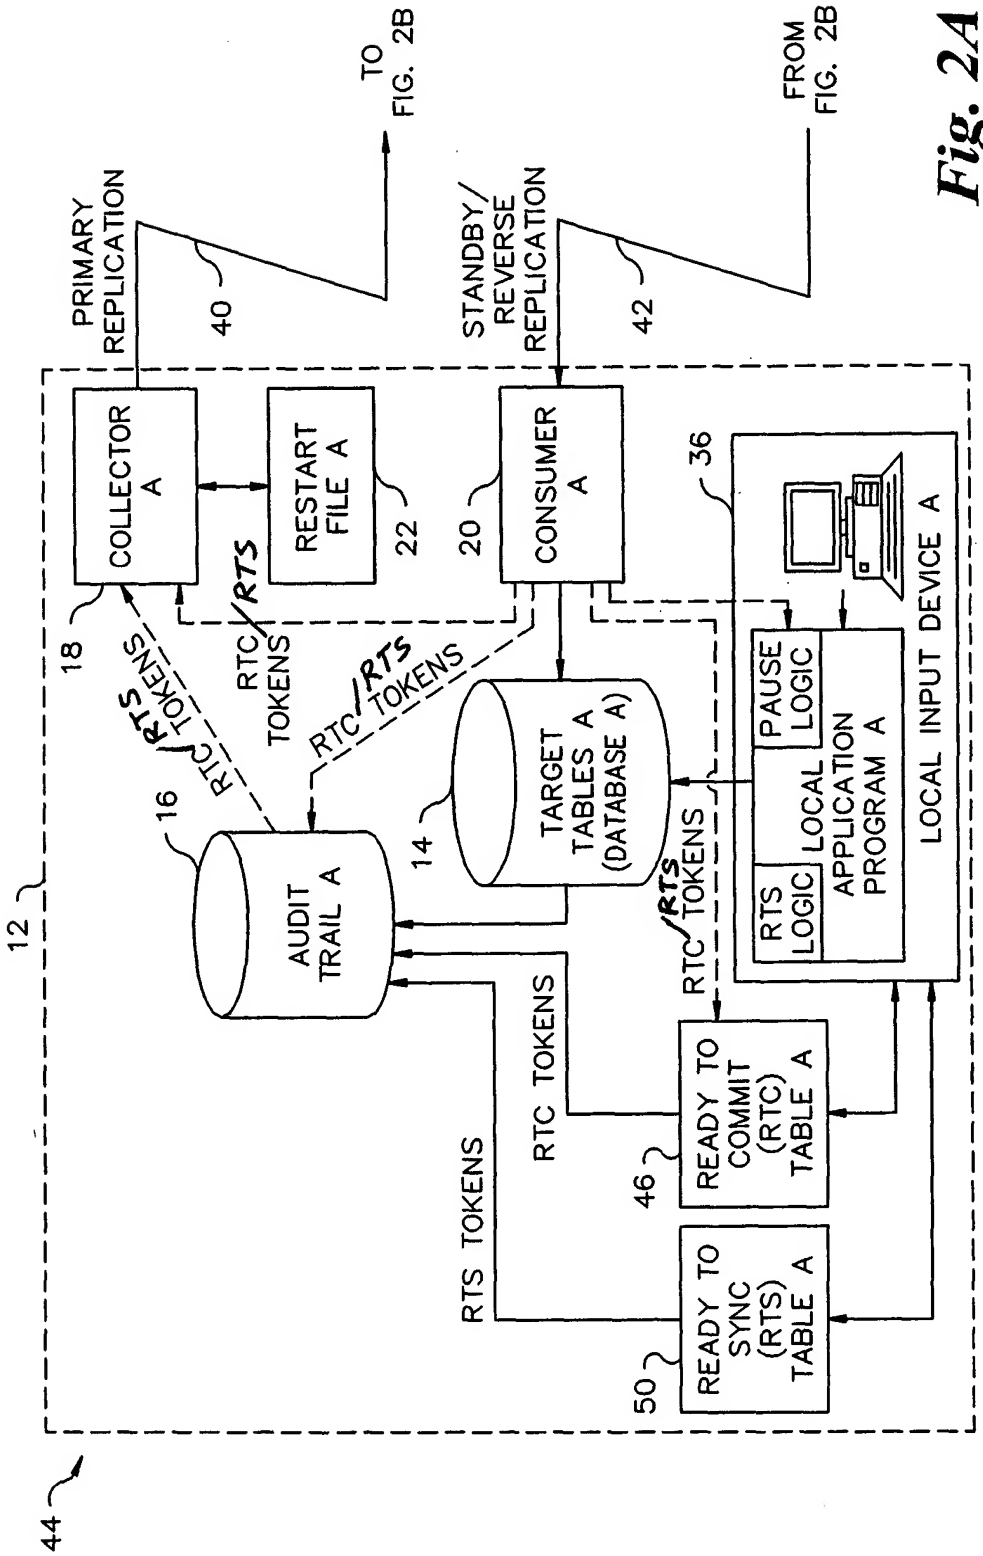
TO FIG. 2B
FROM FIG. 2B
PRIMARY REPLICATION
STANDBY/ REVERSE REPLICATION
40
42
COLLECTOR A
RESTART FILE A
CONSUMER A
18
22
20
36
RTS TOKENS
RTC TOKENS
RTC TOKENS
RTS TOKENS
16
14
AUDIT TRAIL A
TARGET TABLES A (DATABASE A)
PAUSE LOGIC
LOCAL APPLICATION PROGRAM A
RTS LOGIC
LOCAL INPUT DEVICE A
RTC TOKENS
RTS TOKENS
RTC TOKENS
READY TO COMMIT (RTC) TABLE A
READY TO SYNC (RTS) TABLE A
46
50
12
44
NODE A

Fig. 2A

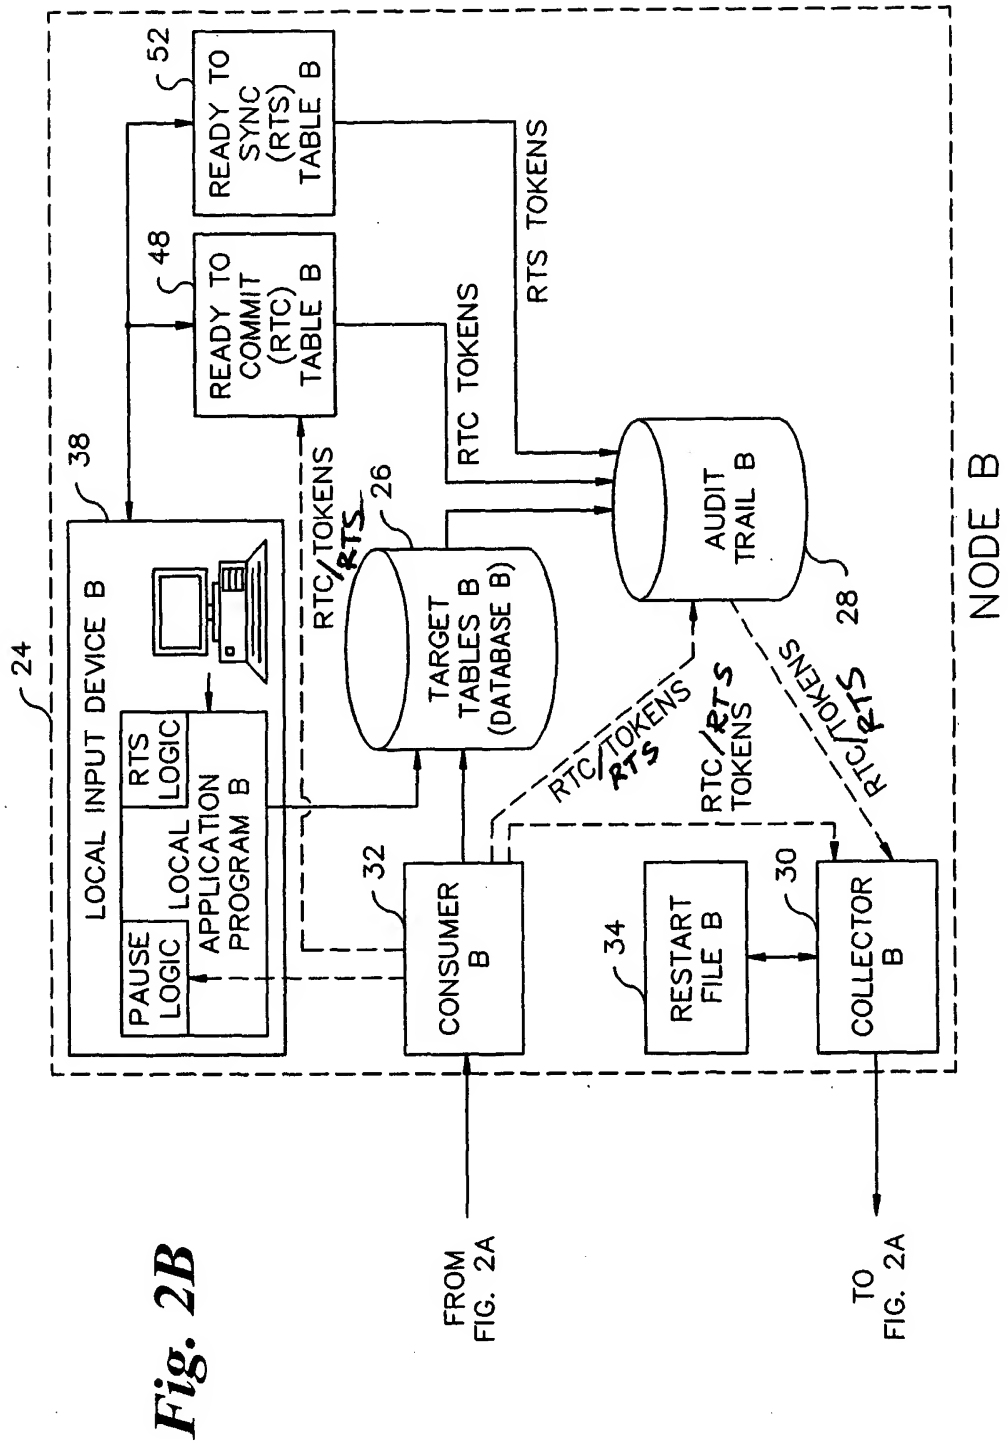
NODE B
LOCAL INPUT DEVICE B
24
38
PAUSE LOGIC
LOCAL APPLICATION PROGRAM B
RTS LOGIC
READY TO SYNC (RTS) TABLE B
52
READY TO COMMIT (RTC) TABLE B
48
RTS TOKENS
RTC TOKENS
RTC/TOKENS
RTS
TARGET TABLES B (DATABASE B)
26
AUDIT TRAIL B
28
CONSUMER B
32
RESTART FILE B
34
COLLECTOR B
30
FROM FIG. 2A
TO FIG. 2A

Fig. 2B

RTC TABLE A

TIME	CONTENTS	
	TRAN ID	FLAG
t_1		
t_2		
t_3		
t_4	101	0
t_5	101	1
t_6	*	
•		
•		
•		
t_n		

AUDIT TRAIL A

TIME	HEADER		DATA
	TRAN ID	TABLE	
t_1	101		BEGIN TRANS. 101
t_2	101	ACCOUNTS	SMITH, JOHN, DEBIT $10
t_3	101	ACCOUNTS	DOE, JANE, CREDIT $10
t_4	101		RTC TOKEN 101
t_5			
t_6	101		COMMIT TRANS. 101

AUDIT TRAIL B

TIME	HEADER		DATA
	TRAN ID	TABLE	
$t_1+\alpha$	101		BEGIN TRANS. 101
$t_2+\alpha$	101	ACCOUNTS	SMITH, JOHN, DEBIT $10
$t_3+\alpha$	101	ACCOUNTS	DOE, JANE, CREDIT $10
$t_4+\alpha$	101		RTC TOKEN 101
$t_5+\alpha$			
$t_6+\alpha$	101		COMMIT TRANS. 101

* NO ENTRY (TRAN ID 101 HAS BEEN DELETED FROM TABLE)

Fig. 3

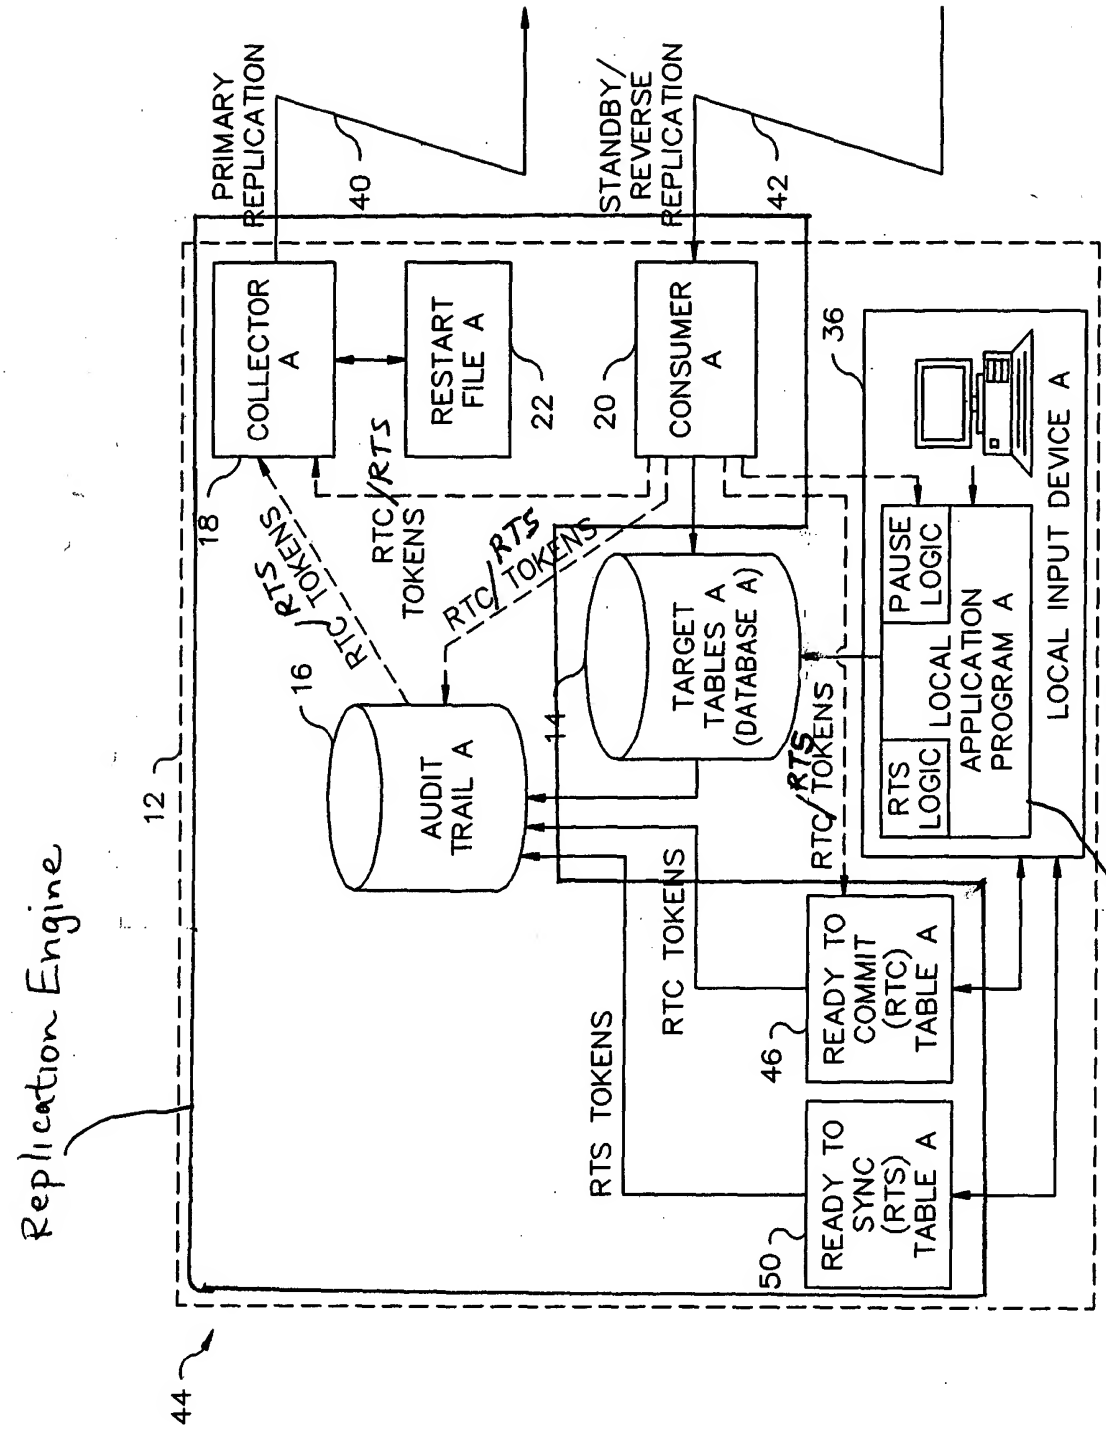
Replication Engine
12
44
PRIMARY REPLICATION
40
STANDBY/ REVERSE REPLICATION
42
COLLECTOR A
18
RESTART FILE A
22
CONSUMER A
20
36
RTS TOKENS
RTC TOKENS
RTC/RTS TOKENS
RTC/RTS TOKENS
16
AUDIT TRAIL A
14
TARGET TABLES A (DATABASE A)
PAUSE LOGIC
LOCAL APPLICATION PROGRAM A
RTS LOGIC
LOCAL INPUT DEVICE A
NODE A
RTC TOKENS
RTS TOKENS
RTC TOKENS
46
READY TO COMMIT (RTC) TABLE A
50
READY TO SYNC (RTS) TABLE A
APPLICATION ENGINE

FIG. 4

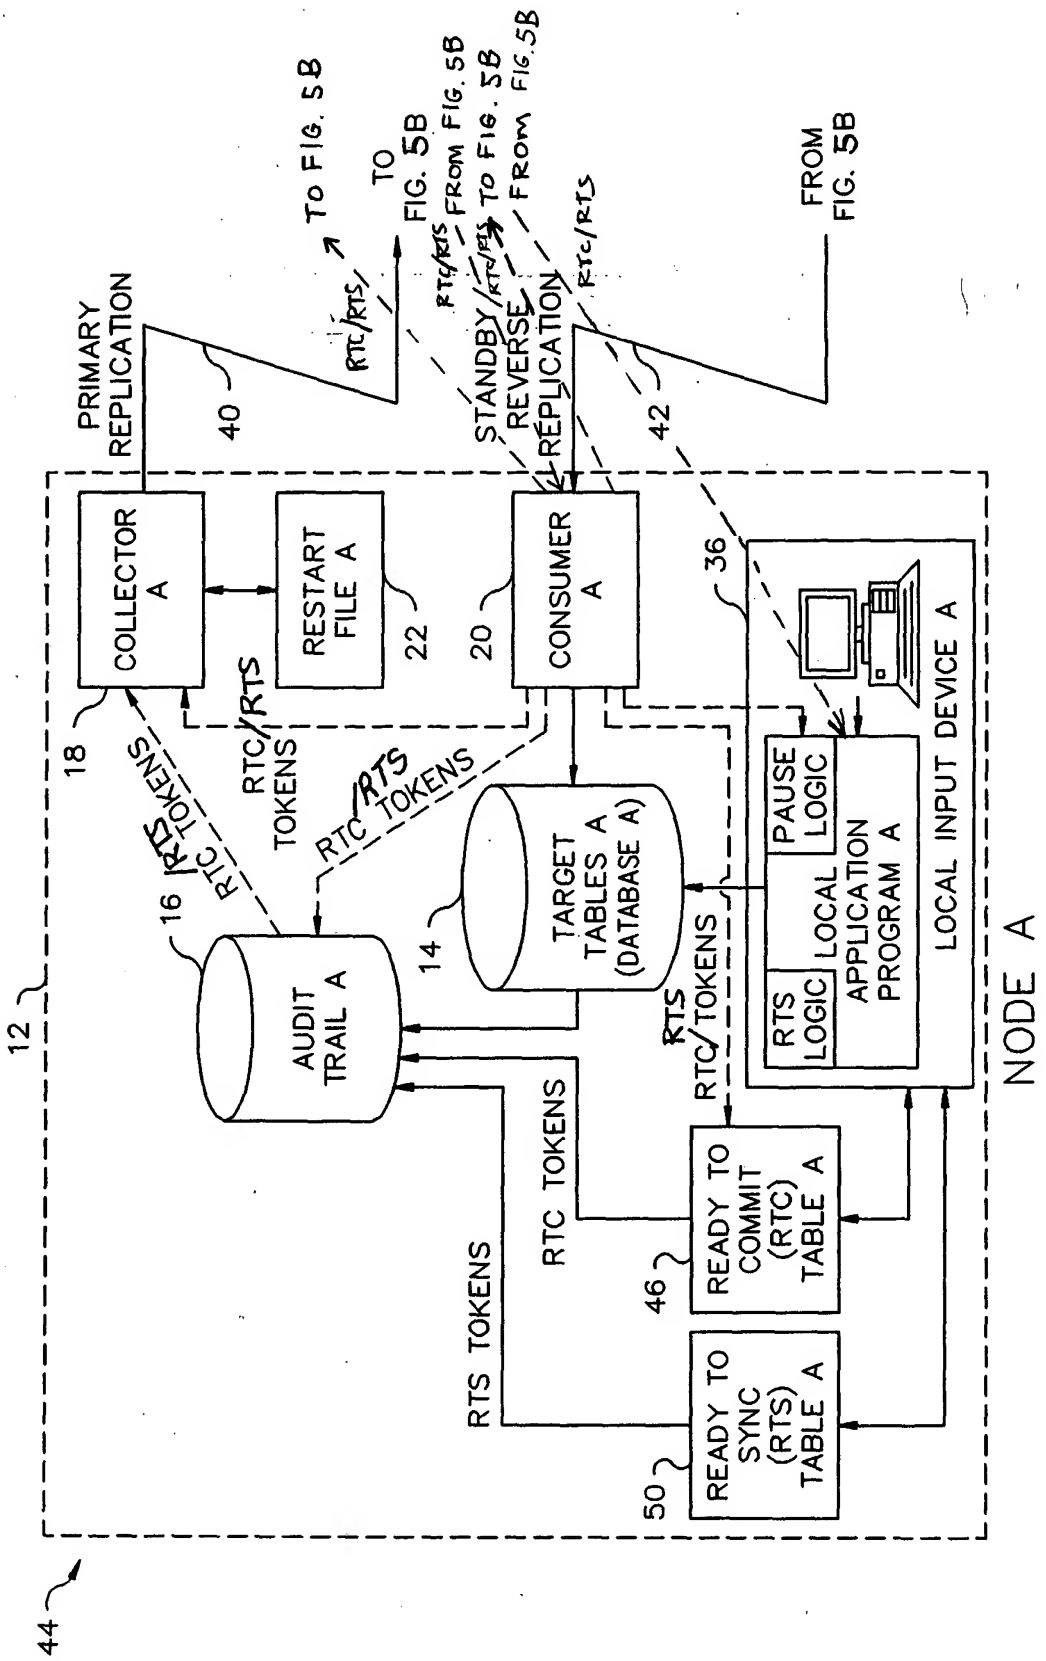
PRIMARY REPLICATION
40
RTC/RTS
TO FIG. 5B
TO FIG. 5B
FROM FIG. 5B
STANDBY REVERSE REPLICATION
RTC/RTS
TO FIG. 5B
FROM FIG. 5B
42
FROM FIG. 5B
COLLECTOR A
18
RESTART FILE A
22
CONSUMER A
20
36
LOCAL INPUT DEVICE A
PAUSE LOGIC
LOCAL APPLICATION PROGRAM A
RTS LOGIC
RTC/RTS TOKENS
RTC TOKENS
RTS TOKENS
16
AUDIT TRAIL A
14
TARGET TABLES A (DATABASE A)
RTS TOKENS
RTC TOKENS
RTS TOKENS
RTC TOKENS
46
READY TO COMMIT (RTC) TABLE A
50
READY TO SYNC (RTS) TABLE A
12
NODE A
44

FIG. 5A

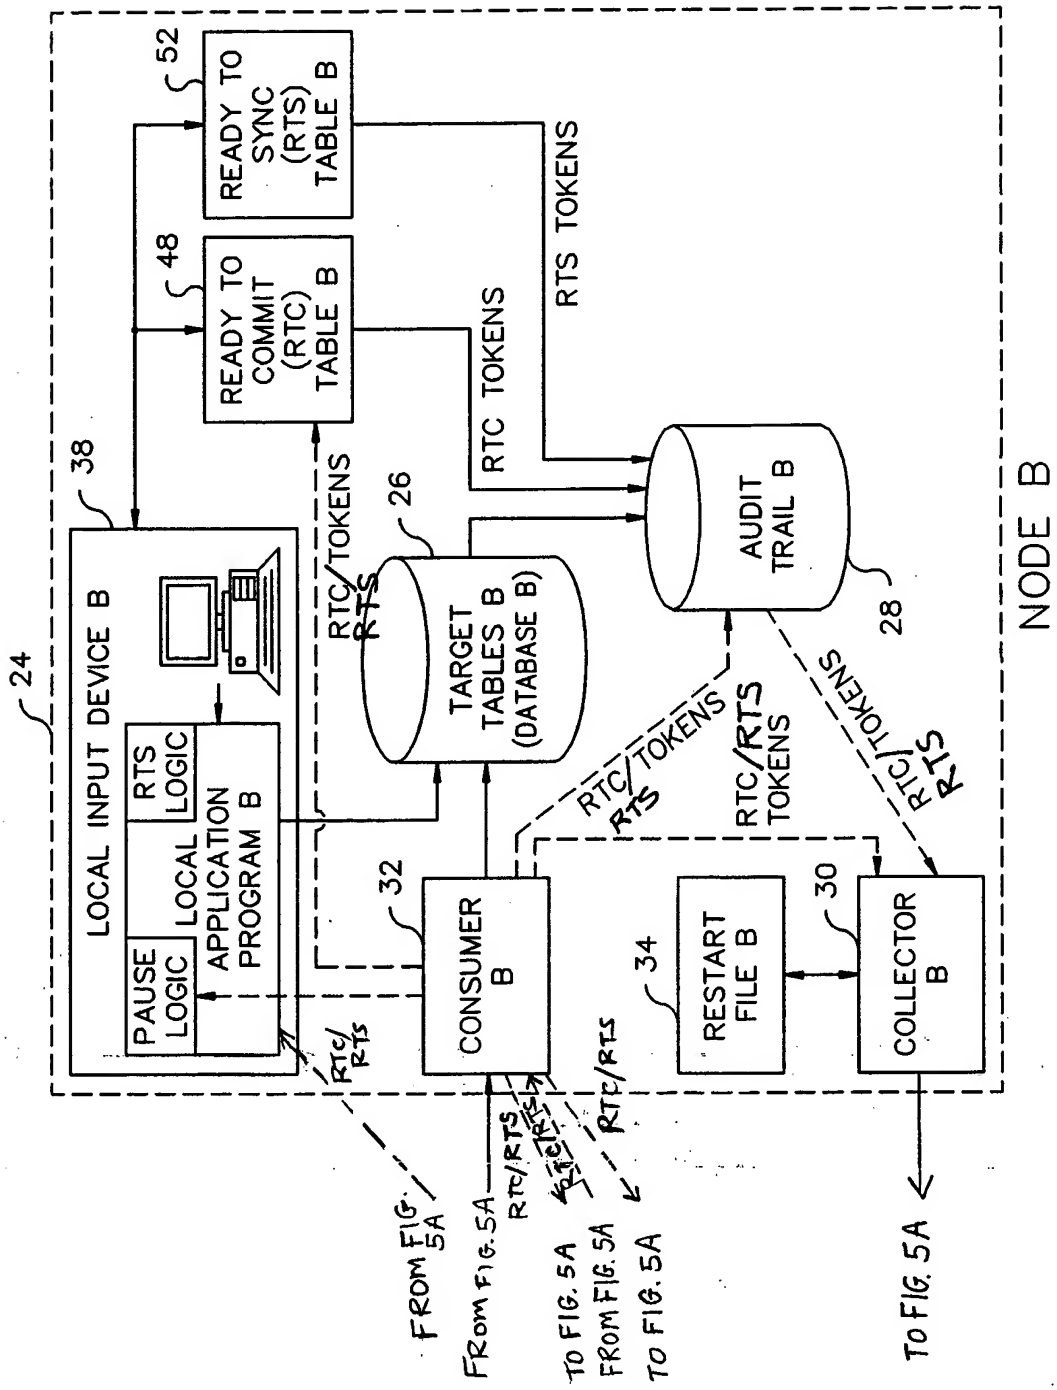
NODE B
LOCAL INPUT DEVICE B
24
38
PAUSE LOGIC
LOCAL APPLICATION PROGRAM B
RTS LOGIC
READY TO SYNC (RTS) TABLE B
52
READY TO COMMIT (RTC) TABLE B
48
RTC TOKENS
RTS TOKENS
RTC/RTS TOKENS
TARGET TABLES B (DATABASE B)
26
AUDIT TRAIL B
28
CONSUMER B
32
RESTART FILE B
34
COLLECTOR B
30
RTC/RTS
FROM FIG. 5A
TO FIG. 5A
FROM FIG. 5A
TO FIG. 5A
TO FIG. 5A

FIG. 5B

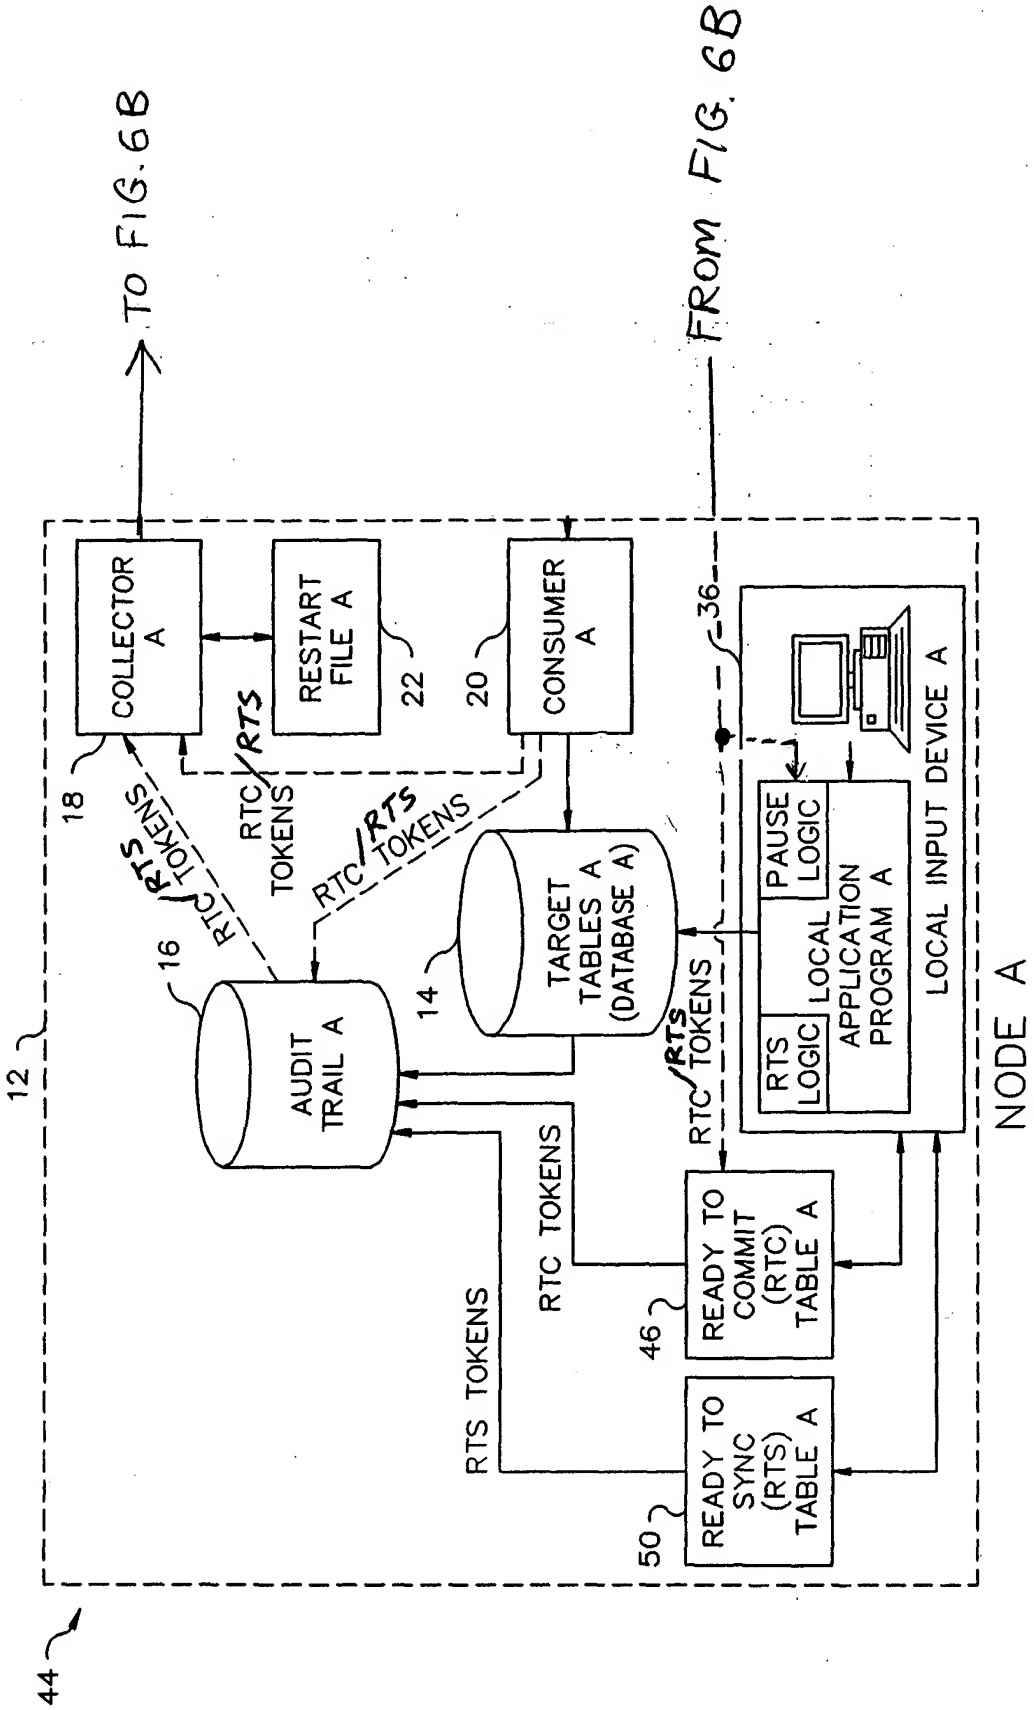
TO FIG. 6B
FROM FIG. 6B
COLLECTOR A
RESTART FILE A
22
CONSUMER A
20
36
18
RTS TOKENS
RTC TOKENS
RTS TOKENS
RTC TOKENS
16
AUDIT TRAIL A
14
TARGET TABLES A (DATABASE A)
PAUSE LOGIC
LOCAL APPLICATION PROGRAM A
RTS LOGIC
LOCAL INPUT DEVICE A
RTC TOKENS
12
NODE A
RTC TOKENS
RTS TOKENS
46
READY TO COMMIT (RTC) TABLE A
50
READY TO SYNC (RTS) TABLE A
44

FIG. 6A

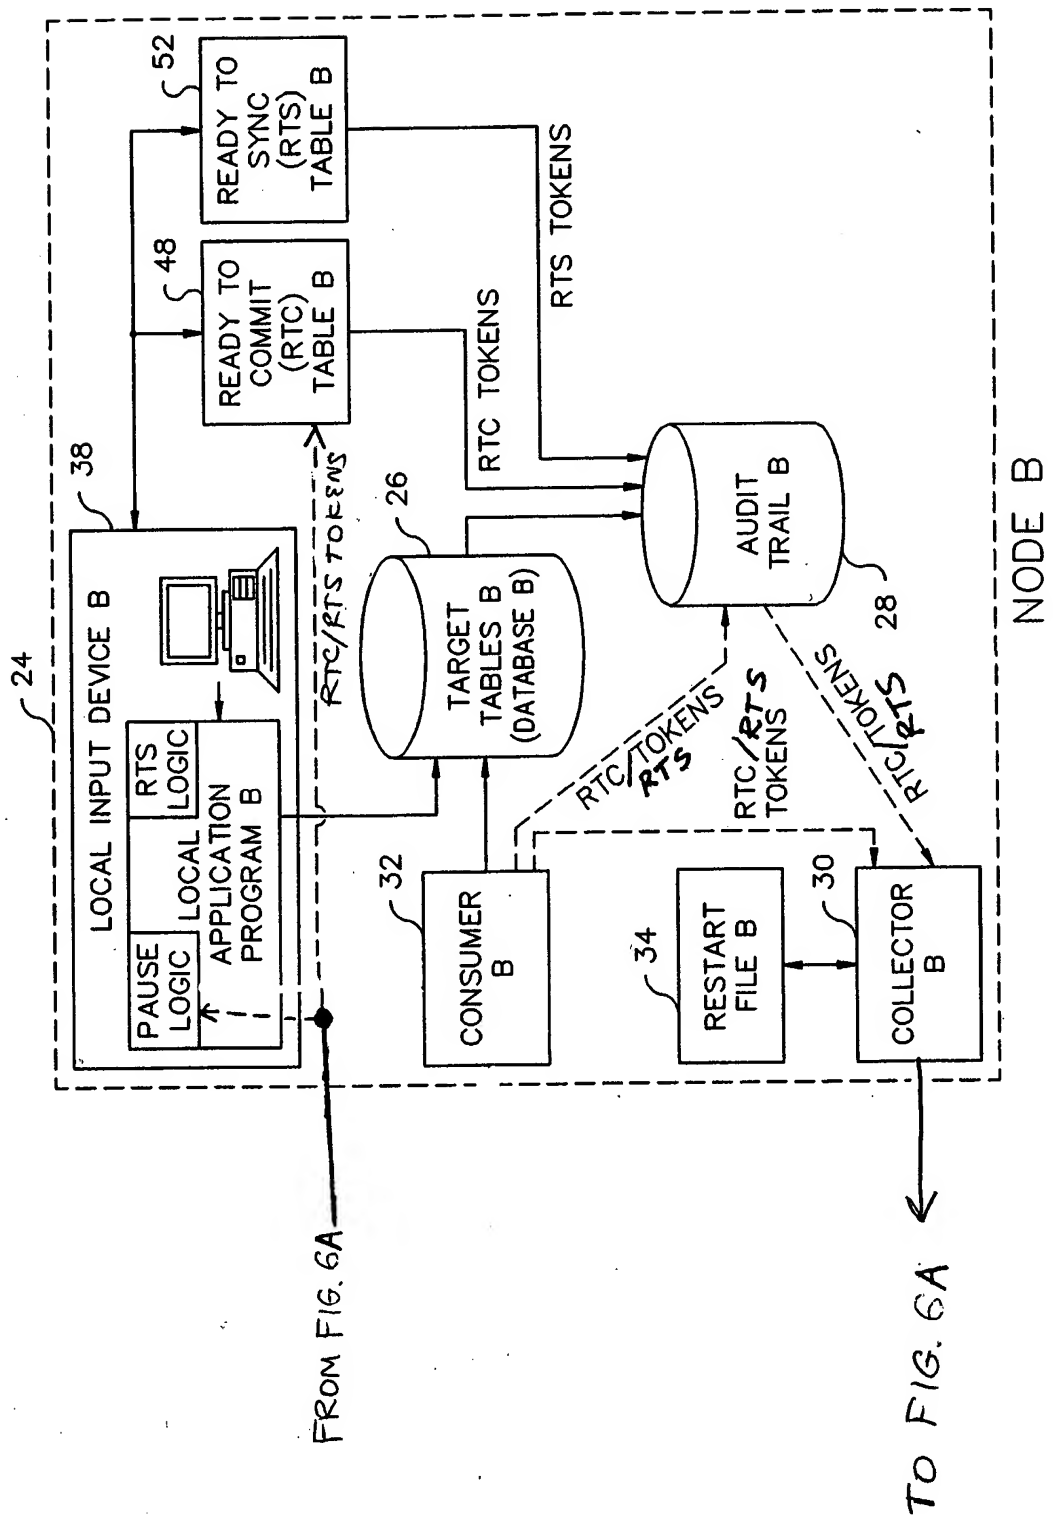
NODE B
24
LOCAL INPUT DEVICE B
38
PAUSE LOGIC
LOCAL APPLICATION PROGRAM B
RTS LOGIC
52
READY TO SYNC (RTS) TABLE B
48
READY TO COMMIT (RTC) TABLE B
RTC/RTS TOKENS
RTS TOKENS
RTC TOKENS
26
TARGET TABLES B (DATABASE B)
AUDIT TRAIL B
28
RTC/RTS TOKENS
RTC/RTS TOKENS
RTC/RTS TOKENS
32
CONSUMER B
34
RESTART FILE B
30
COLLECTOR B
FROM FIG. 6A
TO FIG. 6A

FIG. 6B

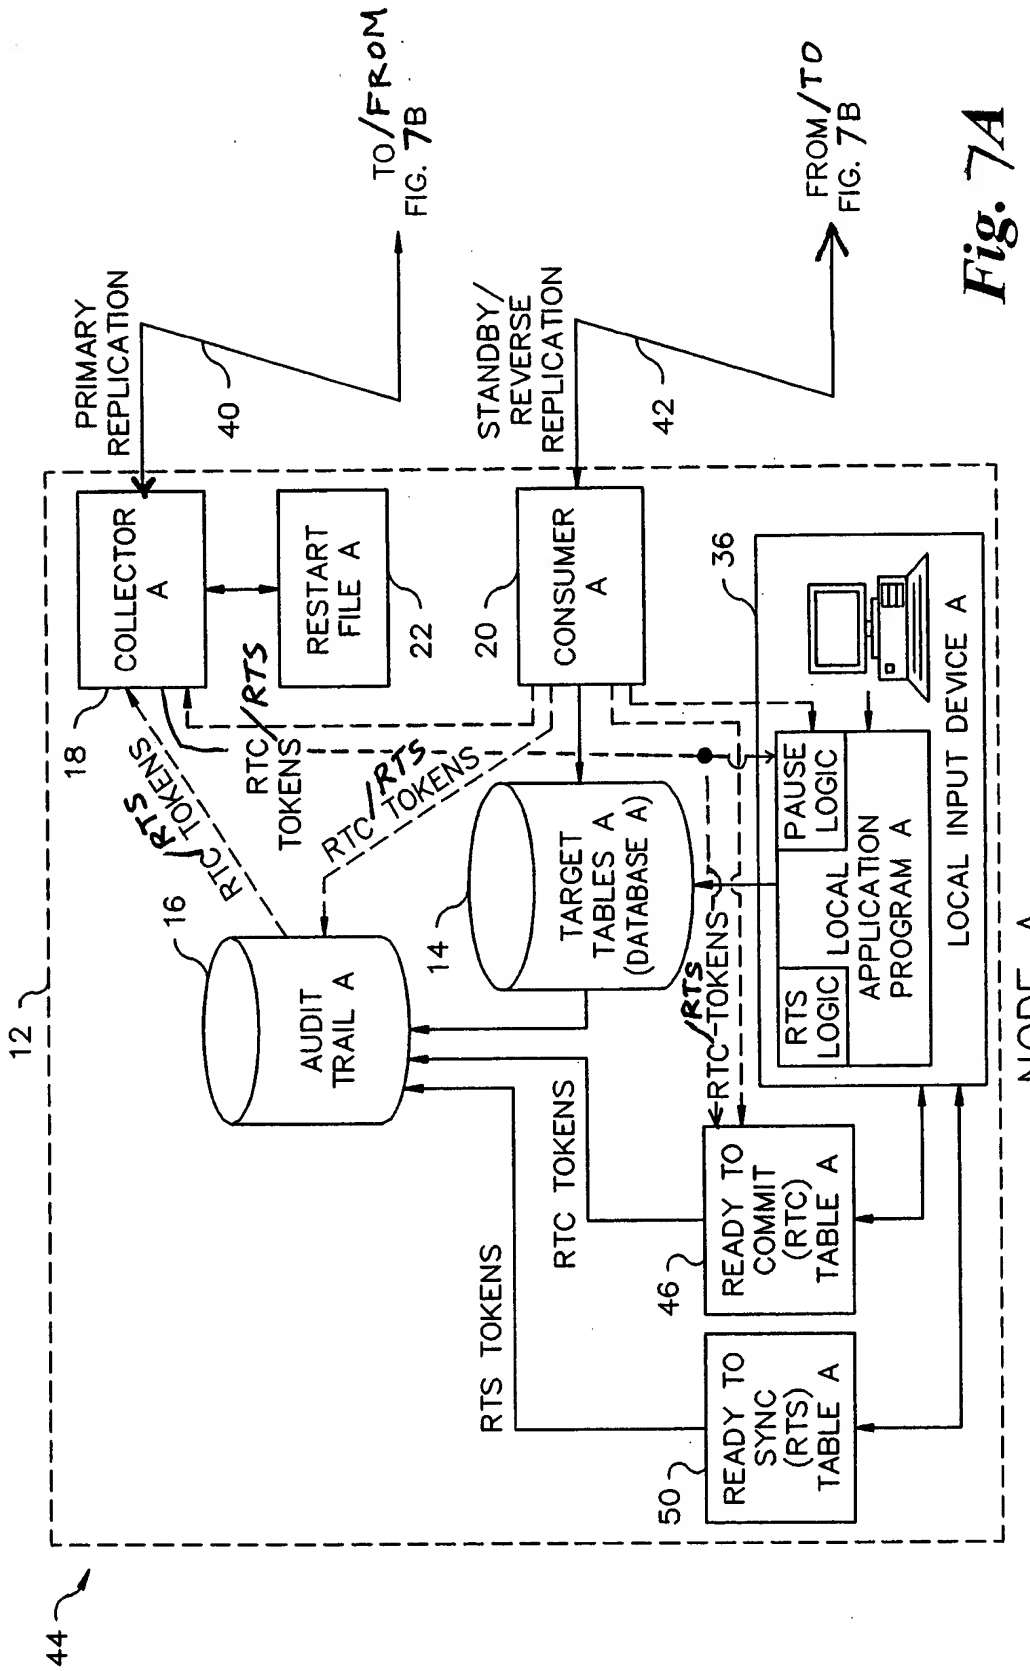
TO/FROM FIG. 7B
FROM/TO FIG. 7B
PRIMARY REPLICATION
STANDBY/ REVERSE REPLICATION
40
42
COLLECTOR A
RESTART FILE A
CONSUMER A
18
22
20
36
RTS TOKENS
RTC TOKENS
RTC/RTS TOKENS
16
14
AUDIT TRAIL A
TARGET TABLES A (DATABASE A)
PAUSE LOGIC
LOCAL APPLICATION PROGRAM A
RTS LOGIC
LOCAL INPUT DEVICE A
RTC/RTS TOKENS
12
RTC TOKENS
RTS TOKENS
READY TO COMMIT (RTC) TABLE A
READY TO SYNC (RTS) TABLE A
46
50
NODE A
44

Fig. 7A

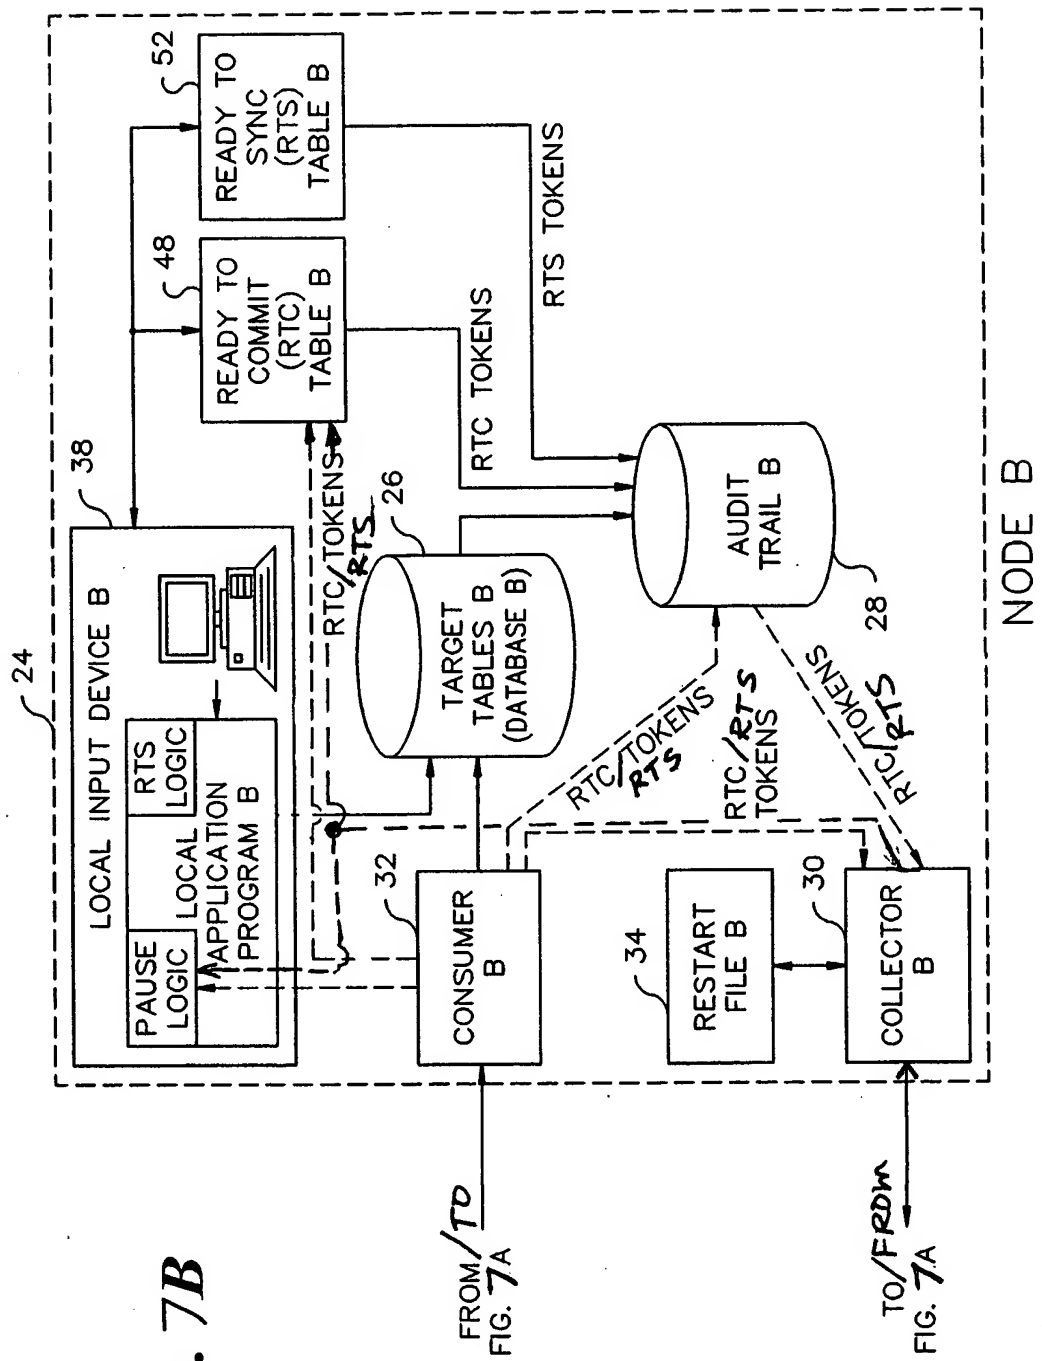
NODE B
LOCAL INPUT DEVICE B
24
38
PAUSE LOGIC
LOCAL APPLICATION PROGRAM B
RTS LOGIC
READY TO SYNC (RTS) TABLE B
52
READY TO COMMIT (RTC) TABLE B
48
RTC TOKENS
RTS TOKENS
RTC/RTS TOKENS
TARGET TABLES B (DATABASE B)
26
AUDIT TRAIL B
28
CONSUMER B
32
RESTART FILE B
34
COLLECTOR B
30
FROM/TO FIG. 7A
TO/FROM FIG. 7A

Fig. 7B

Node	Time	Transaction amount	Comments
A, B	T1		Account balance initially at $1000, both nodes A and B synchronized.
A, B	T2	+$200	Customer deposits (adds) $200 to account balance from node A. Using the synchronous replication method, both nodes are updated to a balance of $1200
A, B	T3	-$75	Customer withdrawals (subtracts) $75 from the account balance from node B. Using the synchronous replication method, both nodes are updated to a balance of $1125
A, B	T4	-$125	Customer transfers (subtracts) $125 from the account balance to a foreign account from node A. Using the synchronous replication method, both nodes are updated to a balance of $1000
	T5		All communications between nodes A and B is lost. A = B = $1000
A	T6	+$100	Customer adds $100 to account balance on node A. A balance = $1100, B balance = $1000
B	T7	-$50	Customer withdrawals $50 from account balance on node B. A node balance = $1100, B node balance = $950.
A, B	T8		Full communications is restored, and the system resolves collisions as follows: A replays the transaction delta changes into B, B replays the transaction delta changes into A. The order of replay could be arbitrary. One could even 'merge' the two replay streams, using timestamps or sequence numbers, etc. The main point is that the account balances will match after all are replayed into the others.
A -> B	T9	+100	More specifically, replaying the A transactions to the B database adds $100 to the B balance of $950, making it $1050. NOTE – we could be done by simply overwriting the A database balance with the B database balance, however replaying in the reverse provides a 'check' of the approach.
B -> A	T10	-50	Replaying the B transactions to the A database subtracts $50 from the A balance of $1100, making it $1050.
A, B	T11		Compare the two balances, A=B and the databases are in sync. If other criteria are met, synchronous replication can be resumed.

Figure 8

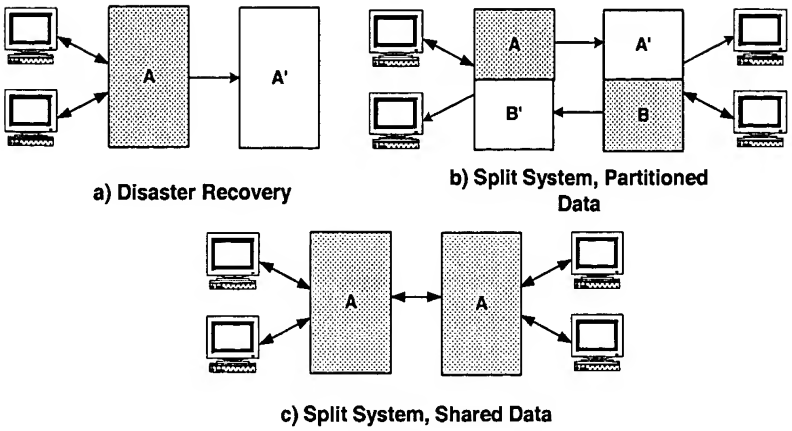

Split System Architectures

Fig. 9

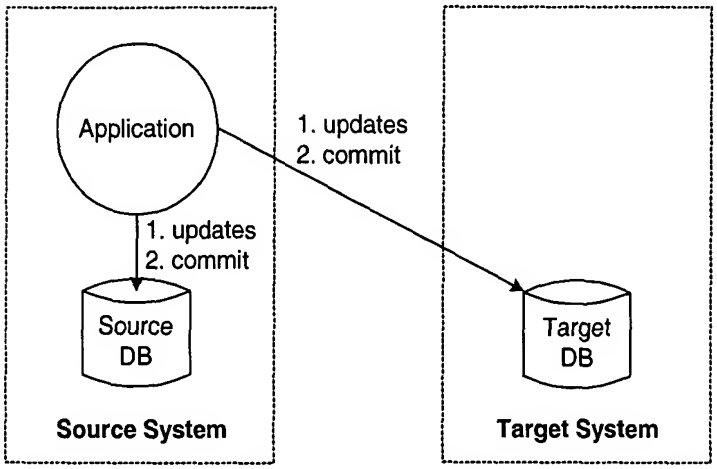

Dual Writes

Fig. 10

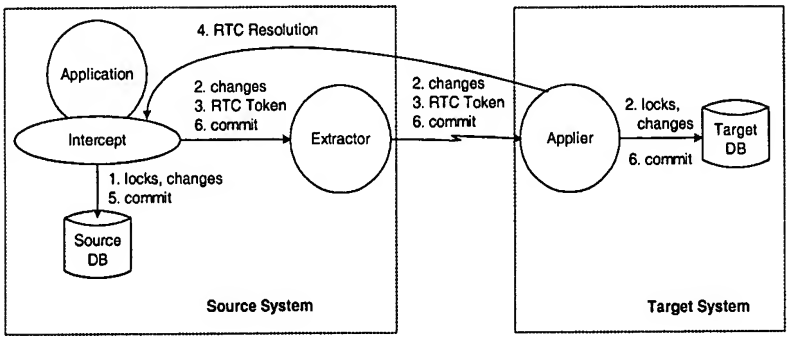

Coordinated Commits

Fig. 11

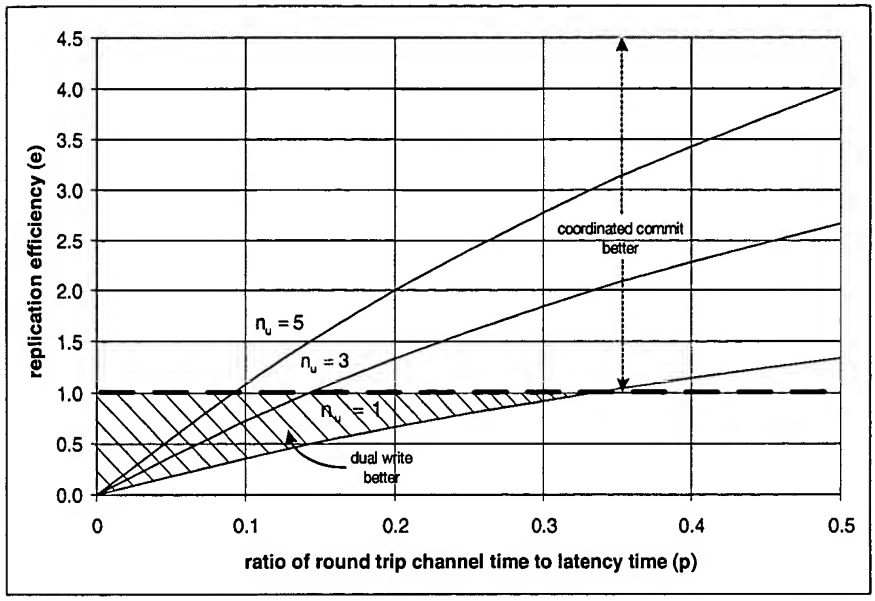

Synchronous Replication Efficiency

Fig. 12

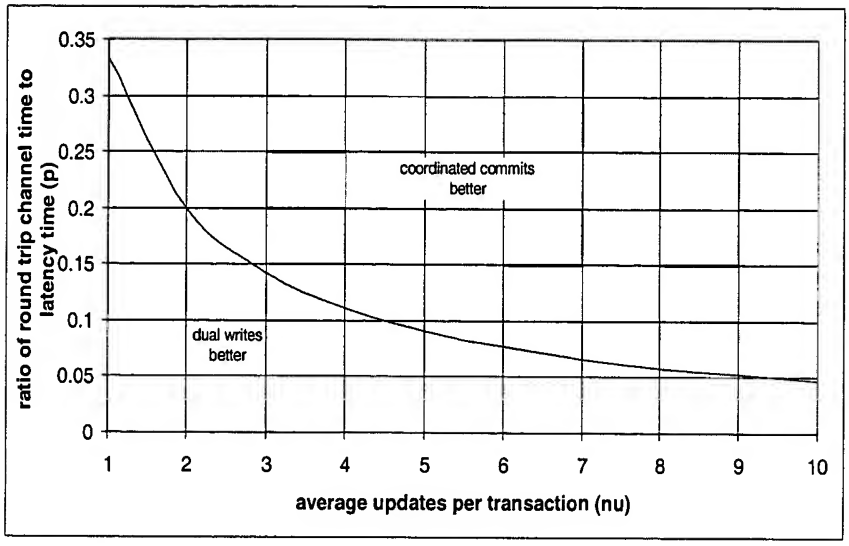

Equal Efficiency (e=1)

Fig. 13